Automating Proofs of Data-Structure Properties in Imperative Programs

Duc-Hiep Chu
National University of Singapore
hiepcd@comp.nus.edu.sg

Joxan Jaffar
National University of Singapore
joxan@comp.nus.edu.sg

Minh-Thai Trinh
National University of Singapore
trinhmt@comp.nus.edu.sg

arXiv:1407.6124v1 [cs.PL] 23 Jul 2014

Abstract

We consider the problem of automated reasoning about dynamically manipulated data structures. The state-of-the-art methods are limited to the unfold-and-match (U+M) paradigm, where predicates are transformed via (un)folding operations induced from their definitions before being treated as *uninterpreted*. However, proof obligations from verifying programs with iterative loops and multiple function calls often do not succumb to this paradigm. Our contribution is a proof method which – beyond U+M – performs *automatic* formula re-writing by treating previously encountered obligations in each proof path as possible *induction hypotheses*. This enables us, for the first time, to systematically reason about a wide range of obligations, arising from practical program verification.

We demonstrate the power of our proof rules on commonly used *lemmas*, thereby close the remaining gaps in existing state-of-the-art systems. Another impact, probably more important, is that our method regains the power of compositional reasoning, and shows that the usage of user-provided lemmas is no longer needed for the existing set of benchmarks. This not only removes the burden of coming up with the appropriate lemmas, but also significantly boosts up the verification process, since lemma applications, coupled with unfolding, often induce very large search space.

1. Introduction

We consider the automated verification of imperative programs with emphasis on reasoning about the functional correctness of dynamically manipulated data structures. In this problem domain, pre/post conditions are specified for each function and an invariant is given for each loop before the reasoning system automatically checks if the program code is correct wrt. the given annotations. The dynamically modified heap poses a big challenge for logical methods. This is because typical correctness properties often require complex combinations of structure, data, and *separation*.

Automated proofs of data structure properties — usually formalized using Separation Logic (or the alike) and extended with *user-defined* recursive predicates — "rely on decidable sub-classes together with the corresponding proof systems based on (un)folding strategies for recursive definitions" [25]. Informally, in the regard of handling recursive predicates, the state-of-the-art [9, 23, 30, 29, 28], collectively called unfold-and-match (U+M) paradigm, employ the basic but systematic transformation steps of *folding* and *unfolding* the rules.

A proof, using U+M, succeeds when we find successive applications of these transformation steps that produce a final formula which is *obviously* provable. This usually means that either (1) there is no recursive predicate in the RHS of the proof obligation and a direct proof can be achieved by consulting some generic SMT solver; or (2) no special consideration is needed on any occurrence of a predicate appearing in the final formula. For example, if $\mathtt{p}(\tilde{u}) \ \wedge \cdots \models \mathtt{p}(\tilde{v})$ is the formula, then this is obviously provable if $\tilde{u}$ and $\tilde{v}$ were *unifiable* (under an appropriate theory governing the meaning of the expressions $\tilde{u}$ and $\tilde{v}$). In other words, we have performed "formula abstraction" [23] by treating the recursively defined term $\mathtt{p}()$ as *uninterpreted*.

Proving Relationship between Unmatchable Predicates

We say, informally, a proof obligation involves *unmatchable* predicates if there exists a recursively defined predicate in the RHS which cannot be transformed, via folding/unfolding, to one that is unifiable with some predicate in the LHS. It can be seen that, U+M (folding/unfolding together with formula abstraction) *cannot* prove relationship between unmatchable predicates.

Let us now highlight scenarios, which are *ubiquitous* in realistic programs, and often lead to proof obligations involving unmatchable predicates. We first articulate them briefly, and then proceed with more examples.

Recursion Divergence
when the "recursion" in the recursive rules is structurally dissimilar to the program code.

Generalization of Predicate
when the predicate describing a loop invariant or a function needs to be used later to prove a weaker property.

First consider "recursion divergence". This happens often with *iterative* programs and when the predicates are not *unary*, i.e., they relate two or more pointer variables, from which the program code traverse or manipulate the data structure in directions different from the definition.

To illustrate, Fig. 1 shows the implementation of a queue using list segments, extracted from the open source program `OpenBSD/queue.h`. Two operations of interest: (1) adding a new element into the end of a non-empty queue (`enqueue`, Fig. 1(a)); (2) deleting an element at the beginning of a non-empty queue (`dequeue`, Fig. 1(b)). A simple property we want to prove is that given a list segment representing a non-empty queue at the beginning, after each operation, we still get back a list segment.

In the two use cases, the "moving pointers" are necessary to recurse differently: the `tail` is moved in `enqueue` while the `head` is moved in `dequeue`. Consequently, no matter how we define list

[Copyright notice will appear here once 'preprint' option is removed.]

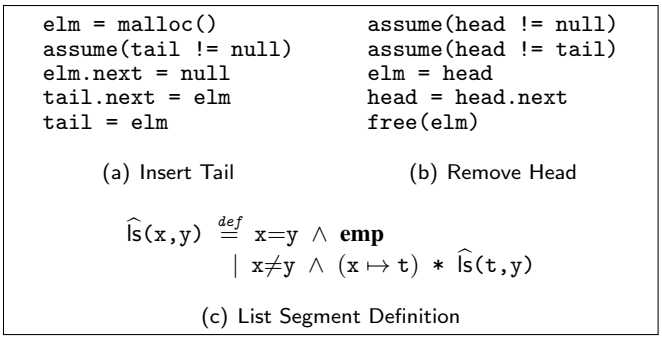

Figure 1: Implementation of a Queue

segments[1], where `head` and `tail` are the two pointers, at least one use case would recurse differently from the definition, thus exhibit the "recursion divergence" scenario and lead to a proof obligation involving unmatchable predicates. More concretely, if list segment is defined as in Fig. 1(c), the `enqueue` operation would lead to an obligation that is impossible for U+M to prove.

Next, we move along to "generalization of predicate". This happens in almost all realistic programs. The reason is because verification of functional correctness is performed *modularly*. More specifically, given the specifications for functions and invariants for loops, we can first perform *local* reasoning before composing the whole proof for the program using, in the context of Separation Logic, the *frame rule* [34]. It can be seen that, given such divide-and-conquer strategy, at the *boundaries* between local code fragments, we would need "generalization of predicate". A particularly important relationship between predicates, at the boundary point, is simply that one (the consequent) is *more general* than the other (the antecedent), representing a valid abstraction step.

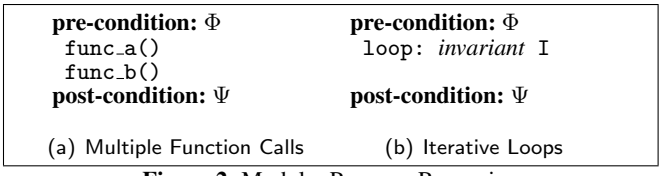

Figure 2: Modular Program Reasoning

Consider the boundaries between function calls, illustrated by the pattern in Fig. 2(a). We start with the pre-condition Φ, calling function `func_a` and then `func_b`. We then need to establish the post-condition Ψ. In traditional forward reasoning, we will write local (and consistent) specifications for `func_a` and `func_b` such that: (1) Φ is stronger than the pre-condition of `func_a`; (2) the post-condition of `func_a` is stronger than the pre-condition of `func_b`; (3) the post-condition of `func_b` is stronger than Ψ. It is hard, if not impossible, to ensure that for each pair (out of three) identified above, the antecedent and the consequent are constructed from matchable predicates. As a concrete example, in `bubblesort` program [9], a boundary between two function calls requires us to prove that a *sorted* linked-list is also a linked-list.

We further argue that in software development, code reuse is often desired. The specification of a function, especially when it is a *library* function, should (or must) be relatively *independent* of the context where the function is plugged in. In each context, we might want to establish arbitrarily different properties, as long as they are *weaker* than what the function can guarantee. In such cases, it is almost certain that we will have proof obligations involving unmatchable predicates.

[1] Typically, list segment can be defined in two ways: the moving pointer is either the *left* one or the *right* one.

Now consider the boundaries caused by loops. In *iterative* algorithms, the loop invariants must be consistent with the code, and yet these invariants are only used later to prove a property often *not* specified using the identical predicates of the invariants. In the pattern shown by Fig. 2(b), this means that the proof obligations relating the pre-condition Φ to the invariant `I` and `I` to post-condition Ψ often involve unmatchable predicates. For example, programs manipulate lists usually have loops of which the invariants need to talk about list segments. Assume that (acyclic) linked-list is defined as below:

$$\mathsf{list}(x) \stackrel{def}{=} x{=}\mathtt{null} \wedge \mathbf{emp} \mid (x \mapsto t) * \mathsf{list}(t)$$

Though $\widehat{\mathsf{ls}}$ and list are closely related, U+M can prove neither of the following obligations:

$$\widehat{\mathsf{ls}}(x, \mathtt{null}) \models \mathsf{list}(x) \tag{1.1}$$

$$\widehat{\mathsf{ls}}(x, y) * \mathsf{list}(y) \models \mathsf{list}(x) \tag{1.2}$$

In summary, the above discussion connects to a serious issue in software development and verification: without the power to relate predicates — when they are unmatchable — *compositional reasoning* is seriously hampered. As a matter of fact, the state-of-the-art is only effective for automatically proving a small fraction of academic and real-world programs.

On using Axioms and Lemmas

We have argued the necessity of proving relationship between unmatchable predicates. Therefore, there cannot be a significant class of programs that is automatically provable using U+M only. How is that there are in fact existing systems displaying proofs of significant examples?

One reason is that some existing methods, e.g., [11, 25, 26, 29], only allow properties to be constructed from a *pre-defined* set of recursive predicates so that hard-wired rules can then be used to facilitate U+M. For systems that support general user-defined predicates [9, 30], they get around the limitation of U+M via the use, without proof, of additional *user-provided* "lemmas" (the corresponding term used in [30] is "axioms"). The general idea is two-fold:

- these lemmas have proofs, though manual, which are simple;
- these lemmas are general and the number of needed ones is small.

Many lemmas used by [30, 28], the most comprehensive existing systems, do not satisfy the two conditions. Some are not so obvious (therefore not reasonable to accept as proven) and some are specifically tailored for the proofs of the target programs (e.g., the lemma used to prove `delete_iter` method in a binary search tree).

As a matter of fact, it is unacceptable that in order to prove more programs, we *continually* add in more custom lemmas to facilitate the proof system. Motivated by this, i.e., to avoid using unproven lemmas, "Cyclic Proof" [6, 7] has recently emerged as an advanced proof technique in Separation Logic. This approach goes beyond U+M and is able to prove the relationship between some commonly used *shape* predicates, even though they are unmatchable.

In certain benchmarks, we concern only with properties involving the "shape" of the data structures, e.g., the obligations (1.1) and (1.2) above. In this regard, "Cyclic Proof" would be quite effective. In practice, however, what is often needed are *custom* predicates for specific application domains, and these will involve *pure* constraints to capture properties of the *data values* in addition to the shapes. Since generally "Cyclic Proof" does not work in the presence of pure constraints, such obligations are out of its scope.

In short, as a proof method, "Cyclic Proof" has advanced the U+M paradigm, but it is still too restrictive to be applicable for the purpose of program verification. We elaborate more in Section 2.

Our Contributions

In this paper, we propose a general proof method that goes significantly beyond the state-of-the-art, namely U+M and "Cyclic Proof", therefore being able to prove relationships between general predicates of arbitrary arity, even when recursive definitions and the code are structurally dissimilar. Our proof method, in one hand, is built upon U+M and therefore subsumes the power of existing (un)-fold reasoners.

On the other hand, in the process of U+M — *systematically* unfolding with the hope to reduce the original obligation to some obligation(s) provable by simple matching and constraint solving — we look for a closely similar obligation in the history of one proof path, of which the truth would *subsume* or *simplify* the current obligation. When such ancestor, called an induction hypothesis, is found, the current obligation can be re-written into a simpler obligation, hopefully then U+M can be effective. Similar to "Cyclic Proof", the soundness of such re-writing steps are based on the theory of induction [18]. In Section 2, we will demonstrate in details the limitation of "Cyclic Proof" as well as why we can consider it as a special and simple realization of the proof rules presented in this paper.

Specifically, this paper makes the following contributions:

- We close the remaining reasoning *gaps* in existing state-of-the-art systems. In particular, we demonstrate in Section 6 that our proof rules can automatically and efficiently discharge all commonly-used lemmas, where existing proof techniques fail to do so.
- We also demonstrate on a comprehensive set of benchmarks, collected from existing systems, that with our proof method the usage of lemmas can be *eliminated*. The impact of this is twofold. First, it means that for proving practical (but small) programs, the users are now free from the burden of providing custom user-defined lemmas. Second, it significantly boosts up the performance, since lemma applications, coupled with unfolding, often induce very large search space.
- The proposed proof method gets us back the power of compositional reasoning in dealing with user-defined recursive predicates. While we have not been able to identify *precisely* the class where our proof method would be effective[2]; we do believe that its potential impact is huge. One important subclass that we can handle effectively is when both the antecedent and the consequent refer to the same structural shape but the antecedent simply makes a *stronger* statement about the values in the structure (e.g., to prove that a sorted list is also a list, an AVL tree is also a binary search tree, a list consists of all data values 999 is one that has all positive data, etc.).

2. The State-of-the-art and Motivating Examples

In this Section, we use illustrative examples to position our proof method against the state-of-the-art.

Unfold-and-Match (U+M) Paradigm

As stated in Section 1, the dominating technique to manipulate user-defined recursive predicates is to employ the basic transformation steps of folding and unfolding the rules, together with formula abstraction, i.e., the U+M paradigm.

[2] This is as *hard* as identifying the class where an invariant discovery technique guarantees to work.

The main challenge of the U+M paradigm is clearly how to systematically search for such sequences of fold/unfold transformations. We believe recent works [23, 30], we shall call the DRYAD works, have brought the U+M to a new level of automation. The key technical step is to use the *program statements* in order to *guide* the sequence of fold/unfold steps of the recursive rules which define the predicates of interest. For example, assume the definition for list segment $\widehat{\mathsf{ls}}$ in Fig. 1(c) and the code fragment in Fig. 3(a).

$\widehat{\mathsf{ls}}$(x,y)	$\widehat{\mathsf{ls}}$(x,y) $*$ (y $\mapsto$ _)
assume(x != null)	
z = x.next	z = y.next
$\widehat{\mathsf{ls}}$(z,y)	$\widehat{\mathsf{ls}}$(x,z)
(a) Code Fragment 1	(b) Code Fragment 2

Figure 3: U+M with List Segments

Here we want to prove that given $\widehat{\mathsf{ls}}$(x,y) at the beginning, we should have $\widehat{\mathsf{ls}}$(z,y) at the end. Since the code touches the "footprint" of x (second statement), it *directs* the unfolding of the predicate containing x, namely $\widehat{\mathsf{ls}}$(x,y), to expose $\mathtt{x} \neq \mathtt{y} \wedge (\mathtt{x} \mapsto \mathtt{t}) * \widehat{\mathsf{ls}}(\mathtt{t},\mathtt{y})$. The consequent can then be established via a simple *matching* from variable z to t.

Now we consider the code fragment in Fig. 3(b): instead of moving one position away from x, we move one away from y. To be convinced that U+M, however, cannot work, it suffices to see that unfolding/folding of $\widehat{\mathsf{ls}}$ does not change the *second argument* of the predicate $\widehat{\mathsf{ls}}$. Therefore, regardless of the unfolding/folding sequence, the arguments y on the LHS and z on the RHS would maintain and can never be matched satisfactorily.

The example in Fig. 3(b) exhibits the "recursion divergence" scenario mentioned in Section 1 and ultimately is about relating two possible definitions of list segment (recursing either on the left or on the right pointer), which U+M fundamentally cannot handle. We will revisit this example in later Sections.

Cyclic Proof

"Cyclic Proof" [6, 7] method operates via the key observation that when two similar obligations are detected in the same proof path, the latter can be used to subsume the former. The soundness is based on the theory of induction. In other words, "Cyclic Proof" can soundly terminate a proof path if the current obligation can be derived from some ancestor obligation in the history of the proof path, just by variable renaming. Of course, we do need to ensure a progressive measure in order to implement "Cyclic Proof". However, for simplicity, we will omit such details in this discussion.

We can view "Cyclic Proof" as a significant extension of U+M. Other than terminating when all the predicates in the RHS can be matched satisfactorily with some predicates in the LHS, as in U+M, it allows termination of a proof path when the *whole* current obligation can be matched with some (whole) obligation appeared in the history of the path. Let us illustrate with a concrete example.

EXAMPLE 1. *Consider the obligation* $\widehat{\mathsf{ls}}(\mathtt{x},\mathtt{y}) * \mathsf{list}(\mathtt{y}) \models \mathsf{list}(\mathtt{x})$, *i.e., obligation* (1.2) *introduced in Section 1.*

See Fig. 4, where the *cyclic path* is at rightmost position. We will focus on this interesting path. The proof proceeds by first unfolding the definition of $\widehat{\mathsf{ls}}$(x,y) in the LHS, followed by unfolding the definition of list(x) in the RHS. The separation rule ($*$) allows us to simplify and subsequently observe the cyclic path, with the appropriate pair denoted with ($\dagger$).

"Cyclic Proof" does open avenue for the integration of user-defined recursive definitions and auxiliary lemmas that relate such definitions, but "it apparently cannot handle definitions where *pure* constraints are present" [38]. In other words, it can only help

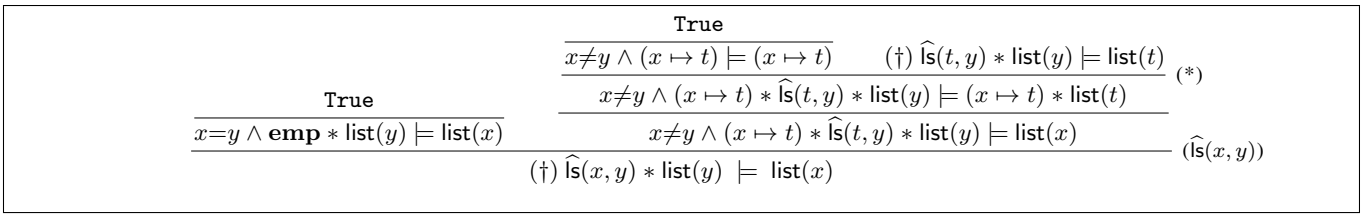
True
$x{=}y \wedge \mathbf{emp} * \mathsf{list}(y) \models \mathsf{list}(x)$
True
$x{\neq}y \wedge (x \mapsto t) \models (x \mapsto t)$
$(\dagger)\ \widehat{\mathsf{ls}}(t,y) * \mathsf{list}(y) \models \mathsf{list}(t)$
$(*)$
$x{\neq}y \wedge (x \mapsto t) * \widehat{\mathsf{ls}}(t,y) * \mathsf{list}(y) \models (x \mapsto t) * \mathsf{list}(t)$
$x{\neq}y \wedge (x \mapsto t) * \widehat{\mathsf{ls}}(t,y) * \mathsf{list}(y) \models \mathsf{list}(x)$
$(\widehat{\mathsf{ls}}(x,y))$
$(\dagger)\ \widehat{\mathsf{ls}}(x,y) * \mathsf{list}(y) \models \mathsf{list}(x)$

Figure 4: Proving with Cyclic Proof

reasoning about the shapes of data structures; but it is not effective when properties of the data values are involved.

Specifically, "Cyclic Proof" is not able to prove the extended versions of obligations (1.1) and (1.2), which are used frequently in [30], when we also want to establish the relationship of collective data values (using sets or sequences) between the LHS and the RHS. It can neither prove that a *sorted* linked-list is also a linked-list (in verifying `bubblesort` program) nor reason about the relationships between *sorted* list segments and *sorted* lists, which are necessary to verify `quicksort` and `mergesort` programs [30].

Also importantly, "Cyclic Proof" cannot handle any obligation of which the LHS and the RHS exhibit *recursion divergence*. For example, it will not be able to prove either $\mathsf{ls(x,y)} * \mathsf{list(y)} \models \mathsf{list(x)}$ or $\widehat{\mathsf{ls}}\mathsf{(x,y)} \models \mathsf{ls(x,y)}$[3], where ls, another form of list segment, is defined as follows.

$$\mathsf{ls(x,y)} \stackrel{def}{=} \mathtt{x}{=}\mathtt{y} \wedge \mathbf{emp} \;\mid\; \mathtt{x}{\neq}\mathtt{y} \wedge (\mathtt{t} \mapsto \mathtt{y}) * \mathsf{ls(x,t)}$$

The reason is that, with recursion divergence, unfolding will typically introduce *new* existential variable(s) on the RHS, which cannot be matched with any variables in the LHS. As a result, (the RHS cannot be simplified and) we will not observe a pair of ancestor-descendant obligations, whose difference is only by variable renaming.

The key distinction that enables us to handle such obligations is that our method treats previous obligations in a proof path as possible (dynamic) induction hypotheses and therefore allows the current obligation not only to be terminated, but also to be *re-written* into a new obligation, usually simpler. The soundness of each rewriting step is also based on the theory of induction. In Section 4.3, we demonstrate our method on these two obligations.

Since "Cyclic Proof" will terminate a proof path and declare the current obligation as proven when it finds a "look-alike" ancestor obligation — achieved by renaming variables on the current obligation, therefore, it can be considered a *special* and simple realization of our proof rules. One can view "Cyclic Proof" to our proof method analogously as cycle detection in explicit model checking to the process of widening[4] in order to discover an invariant with the hope to terminate the reasoning of an unbounded loop.

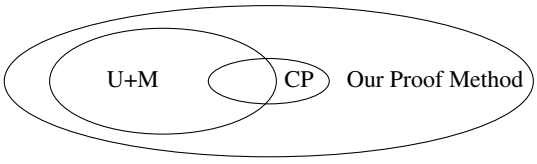

Figure 5: Correlation between U+M, CP, and our proof method.

[3] Such obligations arise often in real programs dealing with lists (e.g., the library **glib/gslist.c**).

[4] Technically, in the context of proving, our re-writing step is a narrowing operation.

To summarize, in Fig. 5, we illustrate the correlation between U+M paradigm, "Cyclic Proof" (CP) method, and our proof method. Theoretically, "Cyclic Proof" subsumes U+M techniques, due to the capability of reasoning about unmatchable predicates. However, in its current implementation, "Cyclic Proof" does not support general non-shape properties (as shown in the SMT-COMP 2014 for Separation Logic). Since such properties (e.g., size property) are ubiquitous in practical programs, there is a large portion of proof obligations that U+M can prove but "Cyclic Proof" cannot. For the domain of interest, our proof method, in contrast, subsumes both the U+M paradigm and "Cyclic Proof" technique. We later demonstrate this clearly in our Experimental Evaluation.

Automatic (and Explicit) Induction

In the literature, there have been works on automatic induction [5, 13, 22, 37]. They are concerned with proving a *fixed* hypothesis, say $h(\tilde{x})$, that is, to show that $h()$ holds over all values of the variables $\tilde{x}$. The challenge is to discover and prove $h(\tilde{x}) \implies h(\tilde{x}')$, where expression $\tilde{x}$ is less than the expression $\tilde{x}'$ in some well-founded measure. Furthermore, a *base case* of $h(\tilde{x}_0)$ needs to be proven. Automating this form of induction usually relies on the fact that some subset of $\tilde{x}$ are variables of *inductive types*.

In contrast, our notion of induction hypothesis is completely different. First, we do not require that some variables are of inductive (and well-founded) types. Second, the induction hypotheses are not supplied explicitly. Instead, they are constructed implicitly via the discovery of a valid proof path. This allows much more potential for automating the proof search. Third, and also different from "Cyclic Proof", multiple induction hypotheses can be exploited within a single proof path. Without this, we would not be able to prove $\widehat{\mathsf{ls}}\mathsf{(x,y)} \models \mathsf{ls(x,y)}$.

3. The Assertion Language

The explicit naming of heaps has emerged naturally in several extensions of Separation Logic (SL) as an aid to practical program verification. Reynolds conjectured that referring explicitly to the current heap in specifications would allow better handles on data structures with sharing [34]. Duck et al. [14], in this vein, extends Hoare Logic with explicit heaps. This extension allows for strongest post conditions, and is therefore suitable for "practical program verification" [8] via constraint-based symbolic execution.

In this paper, we start with the existing specification language in [14], which has two notable features: (a) the use of explicit heap variables, and (b) user-defined recursive properties in a wrapper logic language based on recursive rules. The language provides a new level of expressiveness for specifying properties of heap-manipulating programs. We remark that, common specifications written in traditional Separation Logic, can be automatically compiled into this language.

Due to space limit, we will be brief here and refer interested readers to [14] for more details. A *heap* is a *finite partial map* from *positive* integers to integers, i.e. $\mathsf{Heaps} = \mathbb{Z}_+ \rightharpoonup_{\text{fin}} \mathbb{Z}$. Given a heap $h \in \mathsf{Heaps}$ with domain $D = dom(h)$, we sometimes treat h as the set of pairs $\{(p, v) \mid p \in D \wedge v = h(p)\}$. We note that when a pair (p, v) belongs to some heap h, it is necessary that p is

not a null pointer, i.e., $p \neq 0$. The $\mathcal{H}$-language is the first-order language over heaps.

We use $(*)$ and $(\doteq)$ operators to respectively denote heap disjointness and equation. Intuitively, a constraint like $H \doteq H_1 * H_2$ restricts H_1 and H_2 to be disjoint while giving a name H to the conjoined heaps $H_1 * H_2$.

The Assertion Language $CLP(\mathcal{H})$

As in [14], $\mathcal{H}$ is then extended with *user-defined* recursive predicates. We use the framework of *Constraint Logic Programming* (CLP) [17] to inherit its syntax, semantics, and most importantly, its built-in notions of unfolding rules. For the sake of brevity, we just informally explain the language. The following rules constitutes a recursive definition of predicate $\mathsf{list}(x, L)$, which specifies a skeleton *list*.

$$\mathsf{list}(x, L) \ :- \ \ x = 0, \ L \doteq \Omega.$$
$$\mathsf{list}(x, L) \ :- \ \ L \doteq (x \mapsto t) * L_1, \ \mathsf{list}(t, L_1).$$

The *semantics* of a set of rules is traditionally known as the "least model" semantics (LMS). Essentially, this is the set of groundings of the predicates which are true when the rules are read as traditional implications. The rules above dictates that all true groundings of $\mathsf{list}(x, L)$ are such that x is an integer, L is a heap which contains a skeleton list starting from x. More specifically, when the list is empty, the root node is equal to null ($x = 0$), and the heap is empty ($L \doteq \Omega$). Otherwise, we can split the heap L into two disjoint parts: a singleton heap $(x \mapsto t)$ and the remaining heap L_1, where L_1 corresponds to the heap that contains a skeleton list starting from t.

We now provide the definitions for list segments, which will be used in our later examples. Do note the extra explicit heap variable L, in comparison with corresponding definitions in SL.

$$\widehat{\mathsf{ls}}(x, y, L) \ :- \ \ x = y, \ L \doteq \Omega.$$
$$\widehat{\mathsf{ls}}(x, y, L) \ :- \ \ x \neq y, \ L \doteq (x \mapsto t) * L_1, \ \widehat{\mathsf{ls}}(t, y, L_1).$$

$$\mathsf{ls}(x, y, L) \ :- \ \ x = y, \ L \doteq \Omega.$$
$$\mathsf{ls}(x, y, L) \ :- \ \ x \neq y, \ L \doteq (t \mapsto y) * L_1, \ \mathsf{ls}(x, t, L_1).$$

We also emphasize that the main advantage of this language is the possibility of deriving the strongest postcondition along each program path. It is indeed the main contribution of [14]. Specifically, in order to prove the Hoare triple $\{\phi\}S\{\psi\}$ for a loop-free program S, we simply generate strongest postcondition ψ' along each of its straight-line paths and obtain the verification condition $\psi' \models \psi$. Note that the handling of loops can be reduced to this loop-free setting because of user-specified invariants. We put forward that, in all our experiments (Section 6), the verification conditions are generated using the symbolic execution rules of [14].

4. The Proof Method

Background on CLP: This is provided for the convenience of the readers. An *atom* is of the form $p(\tilde{t})$ where p is a user-defined predicate symbol and $\tilde{t}$ is a tuple of $\mathcal{H}$ terms. A *rule* is of the form $A \texttt{:-} \Psi, \tilde{B}$ where the atom A is the *head* of the rule, and the sequence of atoms $\tilde{B}$ and the constraint Ψ constitute the *body* of the rule. A finite set of rules is then used to define a predicate. A *goal* has exactly the same format as the body of a rule. A goal that contains only constraints and no atoms is called *final*.

A *substitution* θ simultaneously replaces each variable in a term or constraint e into some expression, and we write $e\theta$ to denote the result. A *renaming* is a substitution which maps each variable in the expression into a distinct variable. A *grounding* is a substitution which maps each variable into its intended universe of discourse: an integer or a heap, in the case of our CLP($\mathcal{H}$). Where Ψ is a constraint, a grounding of Ψ results in *true* or *false* in the usual way.

A *grounding* θ of an atom $p(\tilde{t})$ is an object of the form $p(\tilde{t}\theta)$ having no variables. A grounding of a goal $\mathcal{G} \equiv (p(\tilde{t}), \Psi)$ is a grounding θ of $p(\tilde{t})$ where $\Psi\theta$ is *true*. We write $[\![\mathcal{G}]\!]$ to denote the set of groundings of $\mathcal{G}$.

Let $\mathcal{G} \equiv (B_1, \cdots, B_n, \Psi)$ and P denote a non-final goal and a set of rules respectively. Let $R \equiv A \texttt{:-} \Psi_1, C_1, \cdots, C_m$ denote a rule in P, written so that none of its variables appear in $\mathcal{G}$. Let the equation $A = B$ be shorthand for the pairwise equation of the corresponding arguments of A and B. A *reduct* of $\mathcal{G}$ using a clause R, denoted $reduct(\mathcal{G}, R)$, is of the form

$(B_1, \cdots, B_{i-1}, C_1, \cdots, C_m, B_{i+1}, \cdots, B_n, B_i = A, \Psi, \Psi_1)$ provided the constraint $B_i = A \wedge \Psi \wedge \Psi_1$ is satisfiable.

A *derivation sequence* for a goal $\mathcal{G}_0$ is a possibly infinite sequence of goals $\mathcal{G}_0, \mathcal{G}_1, \cdots$, where $\mathcal{G}_i,\ i > 0$ is a reduct of $\mathcal{G}_{i-1}$. A *derivation tree* for a goal is defined in the obvious way.

DEFINITION 1 (Unfold). *Given a program P and a goal $\mathcal{G}$:* UNFOLD$(\mathcal{G})$ *is* $\{\mathcal{G}' | \exists R \in P : \mathcal{G}' = reduct(\mathcal{G}, R)\}$. □

Given a goal $\mathcal{L}$ and an atom $p \in \mathcal{L}$, UNFOLD$_p(\mathcal{L})$ denotes the set of formulas transformed from $\mathcal{L}$ by unfolding p.

DEFINITION 2 (Entailment). *An* entailment *is of the form $\mathcal{L} \models \mathcal{R}$, where $\mathcal{L}$ and $\mathcal{R}$ are goals.* □

This paper considers proving the validity of the entailment $\mathcal{L} \models \mathcal{R}$ under a given program P. This entailment means that $lm(P) \models (\mathcal{L} \rightarrow \mathcal{R})$, where $lm(P)$ denotes the "least model" of the program P which defines the recursive predicates — called *assertion* predicates — occurring in $\mathcal{L}$ and $\mathcal{R}$. This is simply the set of all groundings of atoms of the assertion predicates which are *true* in P. The expression $(\mathcal{L} \rightarrow \mathcal{R})$ means that, for each grounding θ of $\mathcal{L}$ and $\mathcal{R}$, $\mathcal{L}\theta$ is in $lm(P)$ implies that so is $\mathcal{R}\theta$.

4.1 Unfold and Match (U+M)

Assume that we start off with $\mathcal{L} \models \mathcal{R}$. If this entailment can be proved directly, by unification and/or consulting an off-the-shelf SMT solver, we say that the entailment is trivial: a *direct proof* is obtained even without considering the "meaning" of the recursively defined predicates (they are treated as *uninterpreted*). When it is not the case — the entailment is non-trivial — a standard approach is to apply unfolding/folding until all the "frontier" become trivial. We note that, in our framework, we perform only unfolding, but now to both the LHS (the antecedent) and the RHS (the consequent) of the entailment. The effect of unfolding the RHS is similar to a folding operation on the LHS. In more details, when direct proof fails, U+M paradigm proceeds in two possible ways:

- First, select a recursive atom $p \in \mathcal{L}$, unfold $\mathcal{L}$ wrt. p and obtain the goals $\mathcal{L}_1, \ldots, \mathcal{L}_n$. The validity of the original entailment can now be obtained by ensuring the validity of *all* the entailments $\mathcal{L}_i \models \mathcal{R}$ $(1 \leq i \leq n)$.
- Second, select a recursive atom $q \in \mathcal{R}$, unfold $\mathcal{R}$ wrt. q and obtain the goals $\mathcal{R}_1, \ldots, \mathcal{R}_m$. The validity of the original entailment can now be obtained by ensuring the validity of *any one of* the entailments $\mathcal{L} \models \mathcal{R}_j$ $(1 \leq j \leq m)$.

So the proof process can proceed recursively either by proving *all* $\mathcal{L}_i \models \mathcal{R}$ or by proving *one* $\mathcal{L} \models \mathcal{R}_j$ for some j. Since the original LHS and RHS usually contain more than one recursive atoms, this proof process naturally triggers a search tree. Termination can be guaranteed by simply bounding the maximum number of left and right unfolds allowed. In practice, the number of recursive atoms used in an entailment is usually small, such tree size is often manageable.

$$\text{(CP)}\ \frac{\texttt{True}}{\tilde{A} \vdash \mathcal{L} \models \mathcal{R}}\quad \mathcal{L} \models_{\texttt{SMT}} \mathcal{R}\text{, where recursive atoms are treated as uninterpreted}$$

$$\text{(SUB)}\ \frac{\tilde{A} \vdash \mathcal{L} \wedge p(\tilde{x}) \models \mathcal{R}\theta}{\tilde{A} \vdash \mathcal{L} \wedge p(\tilde{x}) \models \mathcal{R} \wedge p(\tilde{y})}\quad \text{there exists a substitution } \theta \text{ for existential variables in } \tilde{y} \text{ s.t. } \mathcal{L} \wedge p(\tilde{x}) \models_{\texttt{SMT}} \tilde{x} = \tilde{y}\theta$$

$$\text{(LU+I)}\ \frac{\bigcup_{i=1}^{n}\{\tilde{A} \cup \{\langle \mathcal{L} \models \mathcal{R}; p\rangle\} \vdash \mathcal{L}_i \models \mathcal{R}\}}{\tilde{A} \vdash \mathcal{L} \models \mathcal{R}}\quad \text{Select an atom } p \in \mathcal{L} \text{ and } \text{UNFOLD}_p(\mathcal{L}) = \{\mathcal{L}_1, \ldots, \mathcal{L}_n\}$$

$$\text{(RU)}\ \frac{\tilde{A} \vdash \mathcal{L} \models \mathcal{R}'}{\tilde{A} \vdash \mathcal{L} \models \mathcal{R}}\quad \text{Select an atom } q \in \mathcal{R} \text{ and } \mathcal{R}' \in \text{UNFOLD}_q(\mathcal{R})$$

$$\text{(IA-1)}\ \frac{\tilde{A} \vdash \mathcal{R}'\theta \wedge \mathcal{L}_2 \models \mathcal{R}}{\tilde{A} \vdash p(\tilde{x}) \wedge \mathcal{L}_1 \wedge \mathcal{L}_2 \models \mathcal{R}}\quad \langle p(\tilde{y}) \wedge \mathcal{L}' \models \mathcal{R}'; p(\tilde{y})\rangle \in \tilde{A} \text{ and } \texttt{gen}(p(\tilde{x})) \geq \texttt{kill}(p(\tilde{y})), \text{ there exists a renaming } \theta \text{ s.t. } \tilde{x} = \tilde{y}\theta \text{ and } \mathcal{L}_1 \models_{\texttt{DP}} \mathcal{L}'\theta$$

$$\text{(IA-2)}\ \frac{\tilde{A} \vdash \mathcal{L}_1 \models \mathcal{L}'\theta}{\tilde{A} \vdash p(\tilde{x}) \wedge \mathcal{L}_1 \models \mathcal{R}}\quad \langle p(\tilde{y}) \wedge \mathcal{L}' \models \mathcal{R}'; p(\tilde{y})\rangle \in \tilde{A} \text{ and } \texttt{gen}(p(\tilde{x})) \geq \texttt{kill}(p(\tilde{y})) \text{ and there exists a renaming } \theta \text{ s.t. } \tilde{x} = \tilde{y}\theta \text{ and } \mathcal{R}'\theta \models_{\texttt{DP}} \mathcal{R}$$

Figure 6: General Proof Rules

4.2 Formula Re-writing with Dynamic Induction Hopotheses

In Section 1, we have identified one major weakness that U+M, which is not being able to relate between unmatchable predicates. Such scenario happens in most realistic programs due to the boundaries created by the multiple function calls and/or iterative loops. We now present a formal calculus for the proof of $\mathcal{L} \models \mathcal{R}$ that goes beyond unfold-and-match. The power of our proof framework comes from the key concept: *induction*.

DEFINITION 3 (Proof Obligation). *A* proof obligation *is of the form $\tilde{A} \vdash \mathcal{L} \models \mathcal{R}$ where $\mathcal{L}$ and $\mathcal{R}$ are goals and $\tilde{A}$ is a set of pairs $\langle A; p\rangle$, where A is an* assumed *entailment and p is a* recursive atom. □

The role of proof obligations is to capture the state of the proof process. Each element in $\tilde{A}$ is a pair, of which the first is an entailment A whose truth can be assumed inductively. A acts as a dynamic induction hypothesis and can be used to transform subsequently encountered obligations in the proof path. We will explain this process later in more details. The second is a recursive atom p, to which the application of a left unfold gives rise to the addition of the induction hypothesis A.

Our proof rules – the obligation at the bottom, and its reduced form on top – are presented in Fig. 6. Given $\mathcal{L} \models \mathcal{R}$, our proof shall start with $\emptyset \vdash \mathcal{L} \models \mathcal{R}$, and proceed by repeatedly applying these rules. Each rule operates on a proof obligation. In this process, the proof obligation may be discharged (indicated by `True`); or new proof obligation(s) may be produced. $\mathcal{L} \models_{\texttt{SMT}} \mathcal{R}$ denotes that the validity of $\mathcal{L} \models \mathcal{R}$ can be obtained simply by consulting a generic SMT solver.

- The *substitution* (SUB) rule removes one occurrence of an assertion predicate, say atom $p(\tilde{y})$, appearing in the RHS of a proof obligation. Applying the (SUB) rule repeatedly will ultimately reduce a proof obligation to the form which contains no recursive atoms in the RHS, while at the same time (hopefully) most existential variables on the RHS are eliminated. Then, the *constraint proof* (CP) rule may be attempted by simply treating all remaining recursive atoms (in the LHS) as uninterpreted and by applying the underlying theory solver assumed in the language we use.

The combination of (SUB) and (CP) rules attempts, what we call, a *direct proof*. In principle, it is similar to the process of "matching" in the U+M paradigm. For brevity we then use $\mathcal{L} \models_{\texttt{DP}}$ $\mathcal{L}$ to denote the fact that the validity of $\mathcal{L} \models \mathcal{R}$ can be proved directly using only (SUB) and (CP) rules.

$$\text{(CP)}\ \frac{\texttt{True}}{\text{(RU)}\ \dfrac{\tilde{A} \vdash \mathsf{list}(x, L) \models x = x, L_1 \doteq \Omega, L \doteq L_1{*}L}{\text{(SUB)}\ \dfrac{\tilde{A} \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, x, L_1), L \doteq L_1{*}L}{\tilde{A} \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2}}}$$

Figure 7: Proving with just U+M

- The *left unfold with induction hypothesis* (LU+I) is a key rule. It selects a recursive atom p on the LHS and performs a complete unfold of the LHS wrt. the atom p, producing a new set of proof obligations. The original obligation, while being removed, is added as an assumption to every newly produced proof obligation, opening the door for the later being used as an induction hypothesis. For technical reason needed below, we do not just add the obligation $\mathcal{L} \models \mathcal{R}$ as an assumption, but also need to keep track of the atom p. This is why in the rule we see a pair $\langle \mathcal{L} \models \mathcal{R}; p\rangle$ added into the current set of assumptions $\tilde{A}$.

On the other hand, the *right unfold* (RU) rule selects some recursive atom q and performs an unfold on the RHS of a proof obligation wrt. q. In the proof process, the two unfold rules will be systematically interleaved.

EXAMPLE 2. *Consider the following proof obligation:*
$\tilde{A} \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2$.

In Fig. 7, we show how this proof obligation can be successfully dispensed by applying (SUB), (RU), and (CP) rules in sequence. Note how the (SUB) rule binds the existential variable y to x and simplifies the RHS of the proof obligation.

- The *induction applications*, namely (IA-1) and (IA-2) rules, transform the current obligation by making use of an assumption which has been added by the (LU+I) rule. The two rules, also called the "induction rules" for short, allow us to treat previously encountered obligations as possible induction hypotheses.

Instead of directly proving the current obligation $\mathcal{L} \models \mathcal{R}$, we now proceed by finding $\overline{\mathcal{L}}$ and $\overline{\mathcal{R}}$ such that $\mathcal{L} \models \overline{\mathcal{L}} \models \overline{\mathcal{R}} \models \mathcal{R}$. The key here is to find those candidate goals where the validity of $\overline{\mathcal{L}} \models \overline{\mathcal{R}}$ directly follows from a "similar" assumption A, together with θ to rename all the variables in A to the variables in the current obligation, namely $\mathcal{L} \models \mathcal{R}$. Assumption A is an obligation which has been previously encountered in the proof process, and

$A\theta$ assumed to be true, as an induction hypothesis. Particularly, we choose $\overline{\mathcal{L}}$ and $\overline{\mathcal{R}}$ so we can (easily) find a renaming θ such that $A\theta \implies \overline{\mathcal{L}} \models \overline{\mathcal{R}}$ ($\implies$ denotes logical implication).

To be more deterministic and to prevent us from transforming to obligations harder than the original one, we require that at least one of the remaining two entailments, namely $\mathcal{L} \models \overline{\mathcal{L}}$ and $\overline{\mathcal{R}} \models \mathcal{R}$, is discharged quickly by a direct proof.

In (IA-1) rule, given the current obligation $p(\tilde{x}) \wedge \mathcal{L}_1 \wedge \mathcal{L}_2 \models \mathcal{R}$ and an assumption $A \equiv p(\tilde{y}) \wedge \mathcal{L}' \models \mathcal{R}'$, we choose $p(\tilde{x}) \wedge \mathcal{L}'\theta \wedge \mathcal{L}_2$ to be our $\overline{\mathcal{L}}$ and $\mathcal{R}'\theta \wedge \mathcal{L}_2$ to be our $\overline{\mathcal{R}}$. We can see that the validity of $\overline{\mathcal{L}} \models \overline{\mathcal{R}}$ directly follows from the assumption $A\theta$. One restriction onto the renaming θ, to avoid circular reasoning, is that θ must rename $\tilde{y}$ to $\tilde{x}$ where $p(\tilde{x})$ is an atom which has been generated after $p(\tilde{y})$ had been unfolded. Such fact is indicated by $\texttt{gen}(p(\tilde{x})) \geq \texttt{kill}(p(\tilde{y}))$ in our rule. While $\texttt{gen}(p)$ denotes the timestamp when the recursive atom p is generated during the proof process, $\texttt{kill}(p)$ denotes the timestamp when p is unfolded and removed. Another side condition for this rule is that the validity of $\mathcal{L} \models \overline{\mathcal{L}}$, or equivalently, $\mathcal{L}_1 \models \mathcal{L}'\theta$ is discharged immediately by a direct proof.

In (IA-2) rule, given the current obligation $p(\tilde{x}) \wedge \mathcal{L}_1 \models \mathcal{R}$ and an assumption $A \equiv p(\tilde{y}) \wedge \mathcal{L}' \models \mathcal{R}'$, on the other hand, $p(\tilde{y})\theta \wedge \mathcal{L}'\theta$ serves as our $\overline{\mathcal{L}}$ while $\mathcal{R}'\theta$ serves as our $\overline{\mathcal{R}}$. The validity of $\overline{\mathcal{L}} \models \overline{\mathcal{R}}$ trivially follows from the assumption $A\theta$, namely $p(\tilde{x}) \wedge \mathcal{L}'\theta \models \mathcal{R}'\theta$. As in (IA-1), we also put similar restriction upon the renaming θ. Another side condition we require is that the validity of $\overline{\mathcal{R}} \models \mathcal{R}$ can be discharged immediately by a direct proof. At this point we could see the duality nature of (IA-1) and (IA-2).

Now let us briefly and intuitively explain the restriction upon the renaming θ. Here we make sure that θ renames atom $p(\tilde{y})$ to atom $p(\tilde{x})$, where $p(\tilde{x})$ has been generated after $p(\tilde{y})$ had been unfolded (and removed). This helps to rule out certain potential θ which does not correspond to a number of left unfolds. Such restriction helps ensure *progressiveness* in the proof process before the induction rules can take place. Otherwise, assuming the truth of $A\theta$ in constructing the proof for A might not be valid. This is the reason why for each element of $\tilde{A}$, we not only keep track of the assumption, but also the recursive atom p to which the application of (LU+I) gives rise to the addition of such assumption.

It is important to note that, our framework as it stands, does not require any consideration of a base case, nor any well-founded measure. Instead, we depend on the Least Model Semantics (LMS) of our assertion language and the above-mentioned restrictions on the renaming θ. In other words, we constrain the use of the rules, which is transparent to the user, in order to achieve a well-founded conclusion.

Let us shed some insights with simple and concrete examples. Consider the recursive predicate p, defined as

$$\texttt{p(}x\texttt{) :- p(}x\texttt{).}$$

and the following two proof obligations:

$$\mathsf{p}(x) \models \mathsf{list}(x, L) \tag{4.1}$$

$$\mathsf{list}(x, L) \models \mathsf{p}(x) \tag{4.2}$$

We will now demonstrate that our method can prove (4.1), but not (4.2). We remark that (4.1) holds because under the Least Model Semantics (LMS) the LHS has no model; therefore no refutation can be found regardless of what is in the RHS [5]. On the other hand, (4.2) does not hold because $x = 0$ (and $L \doteq \Omega$) is a model of the LHS, but not a model of the RHS.

[5] i.e., *false* implies anything

$$\frac{(\dagger)\ \mathsf{p}(x) \models \mathsf{list}(x, L)}{(\dagger)\ \mathsf{p}(x) \models \mathsf{list}(x, L)}\ (\mathsf{p}(x))$$

Figure 8: Cyclic Proof for (4.1)

$$\cfrac{\cfrac{\cfrac{\texttt{True}}{\{A\} \vdash \mathsf{list}(x, L) \models \mathsf{list}(x, L)}\ (\text{CP})}{\{A\} \vdash \mathsf{p}(x) \models \mathsf{list}(x, L)}\ (\text{IA-1})}{\emptyset \vdash \mathsf{p}(x) \models \mathsf{list}(x, L)}\ (\text{LU+I})$$

Figure 9: Our Proof for (4.1)

We start the discussion by first showing how "Cyclic Proof" would handle (4.1). See Fig. 8. After unfolding the predicate $\mathsf{p}(x)$ in the LHS, we immediately observe the same obligation, i.e., a cyclic path has been established. Consequently, we can terminate the proof path and discharge the original obligation.

We now proceed showing in Fig. 9 how our method would handle this obligation. We first perform a left unfolding, adding $A \equiv \langle \mathsf{p}(x) \models \mathsf{list}(x, L); \mathsf{p}(x) \rangle$ into the set of assumptions. Note that this unfolding step kills the predicate $\mathsf{p}(x)$ and generates a new predicate $\mathsf{p}(x)$. Thus the rule (IA-1) is applicable now. We then rewrite the LHS from $\mathsf{p}(x)$ to $\mathsf{list}(x, L)$. Finally the proof succeeds by consulting constraint solver, treating $\mathsf{list}(x, L)$ as uninterpreted.

$$\cfrac{\cfrac{\{A'\} \vdash x = 0, L \doteq \Omega \models \mathsf{p}(x) \qquad \vdots}{\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{p}(x)}\ (\text{LU+I})}{\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{p}(x)}\ (\text{RU})$$

Figure 10: An Unsuccessful Attempt for (4.2)

In contrast, now consider obligation (4.2) in Fig 10. Obviously, a direct proof for this is not successful. However, if we proceed by a right unfold first, we get back the same obligation. Different from before, and importantly, now no new assumption is added. We can see that the step does not help us progress and therefore performing right unfold repetitively would get us nowhere. Now consider performing a left unfold on the obligation. The proof succeeds if we can discharge both

$$\{A'\} \vdash x = 0, L \doteq \Omega \models p(x) \ \texttt{and}$$
$$\{A'\} \vdash L \doteq (x \mapsto t) * L_1, \mathsf{list}(t, L_1) \models p(x),$$

where $A' \equiv \langle \mathsf{list}(x, L) \models \mathsf{p}(x); \mathsf{list}(x, L) \rangle$.

Focus on the obligation $\{A'\} \vdash x = 0, L \doteq \Omega \models p(x)$. Clearly consulting a constraint solver or performing substitution does not help. Rule (LU+I) is not applicable since no recursive predicate on the LHS. As before, we cannot progress using (RU) rule. Importantly, the side conditions prevent (IA-1) and (IA-2) from taking place. In summary, with our proof rules, this (wrong fact) cannot be established.

4.3 Proving the Two Motivating Examples

Let us now revisit the two motivating examples introduced earlier, on which both U+M and "Cyclic Proof" are not effective. The main reason is that both examples involve unmatchable predicates while at the same time exhibiting "recursion divergence".

EXAMPLE 3. *Consider the entailment relating two definitions of list segments:* $\widehat{\mathsf{ls}}(x, y, L) \models \mathsf{ls}(x, y, L)$.

Our method can discharge this obligation by applying (IA-1) rule twice. For space reason, in Fig. 11, we only show the interesting path of the proof tree (leftmost position). First, we unfold the predicate $\widehat{\mathsf{ls}}(x, y, L)$ in the LHS of the given obligation via (LU+I) rule. The original obligation, while being removed, is added as an assumption A_1. We next make use of A_1 as an induction hypothesis

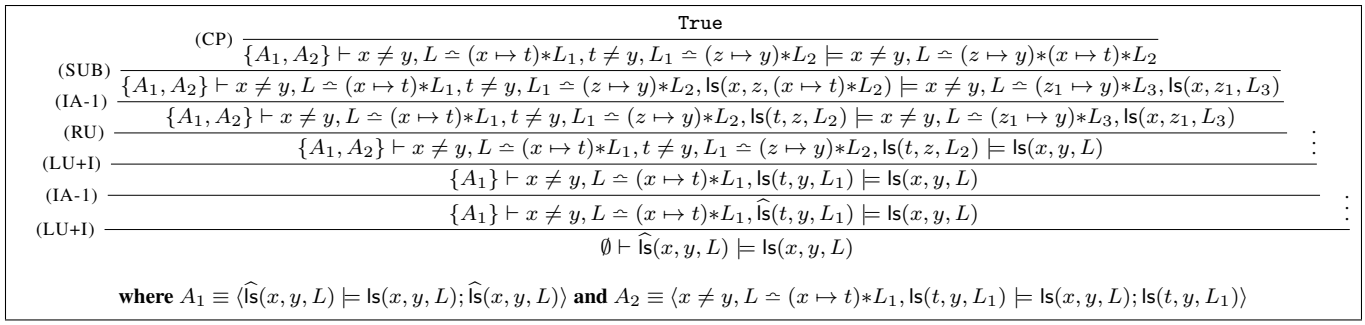
(CP) True
{A1, A2} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, t ≠ y, L1 ≐ (z ↦ y)*L2 ⊨ x ≠ y, L ≐ (z ↦ y)*(x ↦ t)*L2
(SUB) {A1, A2} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, t ≠ y, L1 ≐ (z ↦ y)*L2, ls(x, z, (x ↦ t)*L2) ⊨ x ≠ y, L ≐ (z1 ↦ y)*L3, ls(x, z1, L3)
(IA-1) {A1, A2} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, t ≠ y, L1 ≐ (z ↦ y)*L2, ls(t, z, L2) ⊨ x ≠ y, L ≐ (z1 ↦ y)*L3, ls(x, z1, L3)
(RU) {A1, A2} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, t ≠ y, L1 ≐ (z ↦ y)*L2, ls(t, z, L2) ⊨ ls(x, y, L)
(LU+I) {A1} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, ls(t, y, L1) ⊨ ls(x, y, L)
(IA-1) {A1} ⊢ x ≠ y, L ≐ (x ↦ t)*L1, ls^(t, y, L1) ⊨ ls(x, y, L)
(LU+I) ∅ ⊢ ls^(x, y, L) ⊨ ls(x, y, L)
where A1 ≡ ⟨ls^(x, y, L) ⊨ ls(x, y, L); ls^(x, y, L)⟩ and A2 ≡ ⟨x ≠ y, L ≐ (x ↦ t)*L1, ls(t, y, L1) ⊨ ls(x, y, L); ls(t, y, L1)⟩

Figure 11: Proving $\widehat{\mathsf{ls}}(x, y, L) \models \mathsf{ls}(x, y, L)$.

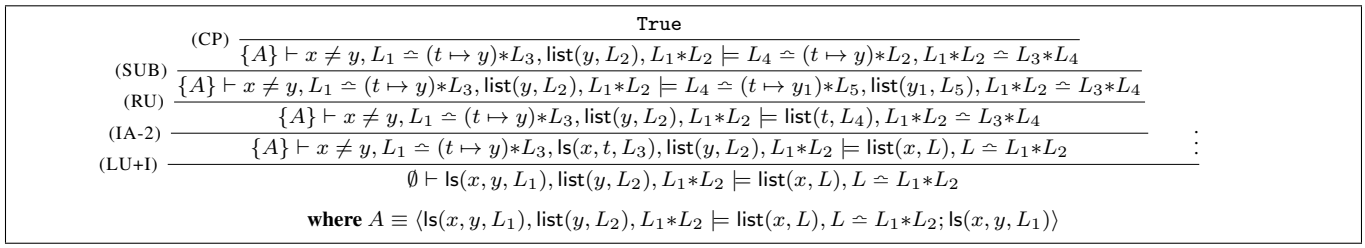
(CP) True
{A} ⊢ x ≠ y, L1 ≐ (t ↦ y)*L3, list(y, L2), L1*L2 ⊨ L4 ≐ (t ↦ y)*L2, L1*L2 ≐ L3*L4
(SUB) {A} ⊢ x ≠ y, L1 ≐ (t ↦ y)*L3, list(y, L2), L1*L2 ⊨ L4 ≐ (t ↦ y1)*L5, list(y1, L5), L1*L2 ≐ L3*L4
(RU) {A} ⊢ x ≠ y, L1 ≐ (t ↦ y)*L3, list(y, L2), L1*L2 ⊨ list(t, L4), L1*L2 ≐ L3*L4
(IA-2) {A} ⊢ x ≠ y, L1 ≐ (t ↦ y)*L3, ls(x, t, L3), list(y, L2), L1*L2 ⊨ list(x, L), L ≐ L1*L2
(LU+I) ∅ ⊢ ls(x, y, L1), list(y, L2), L1*L2 ⊨ list(x, L), L ≐ L1*L2
where A ≡ ⟨ls(x, y, L1), list(y, L2), L1*L2 ⊨ list(x, L), L ≐ L1*L2; ls(x, y, L1)⟩

Figure 12: Proving $\mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L_1{*}L_2 \models \mathsf{list}(x, L), L \doteq L_1{*}L_2$.

to perform a re-writing step, i.e., an application of (IA-1) rule. Similarly, in the third step, we unfold the predicate $\mathsf{ls}(t, y, L_1)$ in the LHS via (LU+I) rule and add the assumption A_2. After unfolding in the RHS via (RU) rule and re-writing with the induction hypothesis A_2 using (IA-1) rule, we are able to bind the existential variable z_1 to z and simplify both sides of the proof obligation using (SUB) rule. Finally, the proof path is terminated by consulting a constraint solver, i.e., (CP) rule.

EXAMPLE 4. *Consider the entailment:*
$\mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L_1{*}L_2 \models \mathsf{list}(x, L), L \doteq L_1{*}L_2$.

Fig. 12 shows, only the interesting proof path, how we can successfully prove this entailment using the (IA-2) rule. We first unfold $\mathsf{ls}(x, y, L_1)$ in the LHS, adding A into the set of assumptions. Then using A as an induction hypothesis, we can rewrite the current obligation via (IA-2) rule. Note that, here we use (IA-2) rule instead of (IA-1) rule as in previous example. After applying (RU) rule, we are able to bind the existential variable y_1 to y and simplify both sides of the proof obligation with (SUB) rule. Finally, the proof path is terminated by consulting a constraint solver, i.e., (CP) rule.

Let us pay a closer attention at the step where we attempt re-writing, making using the available induction hypothesis. For the sake of discussion, assume that instead of (IA-2) we now attempt to apply rule (IA-1). The requirement for θ forces it to rename x to x and y to t. However, the side condition $\mathcal{L}_1 \models_{\mathtt{DP}} \mathcal{L}'\theta$ cannot be fulfilled, since $x \neq y, L_1 \doteq (t \mapsto y){*}L_3, \mathsf{list}(y, L_2), L_1{*}L_2 \not\models_{\mathtt{DP}} \mathsf{list}(t, \_)$.

Now return to the attempt of (IA-2) rule. The RHS of the current obligation matches with the RHS of the *only* induction hypothesis perfectly. This matching requires θ to rename x back to x. On the LHS, we further require θ to rename y to t so that $\mathsf{ls}(x, t) \equiv \mathsf{ls}(x, y)\theta$. Note that $\mathsf{ls}(x, t)$ was indeed generated after $\mathsf{ls}(x, y)$ had been unfolded and removed (i.e., killed). The remaining transformation is straightforward.

4.4 Soundness

THEOREM 1 (Soundness). *An entailment $\mathcal{L} \models \mathcal{R}$ holds if, starting with $\emptyset \vdash \mathcal{L} \models \mathcal{R}$, there exists a sequence of applications of proof rules that results in an empty set of proof obligations.* □

Proof Sketch. The soundness of rule (CP) is obvious. The rule (RU) is sound because when $\mathcal{R}' \in \text{UNFOLD}_q(\mathcal{R})$ then $\mathcal{R}' \models \mathcal{R}$. Therefore, the proof of $\tilde{A} \vdash \mathcal{L} \models \mathcal{R}$ can be replaced by the proof of $\tilde{A} \vdash \mathcal{L} \models \mathcal{R}'$ since $\mathcal{L} \models \mathcal{R}'$ is stronger than $\mathcal{L} \models \mathcal{R}$. Similarly, the rule (SUB) is sound because $\mathcal{L} \wedge p(\tilde{x}) \models \mathcal{R}\theta$ and $\mathcal{L} \wedge p(\tilde{x}) \models_{\mathtt{CP}} \tilde{x} = \tilde{y}\theta$ is stronger than the $\mathcal{L} \wedge p(\tilde{x}) \models \mathcal{R} \wedge p(\tilde{y})$.

The rule (LU+I) is *partially sound* in the sense that when $\text{UNFOLD}_p(\mathcal{L}) = \{\mathcal{L}_1, \ldots, \mathcal{L}_n\}$, then proving $\mathcal{L} \models \mathcal{R}$ can be substituted by proving $\mathcal{L}_1 \models \mathcal{R}, \ldots, \mathcal{L}_n \models \mathcal{R}$. This is because in the least-model semantics of the definitions, $\mathcal{L}$ is equivalent to $\mathcal{L}_1 \vee \ldots \vee \mathcal{L}_n$. However, whether the addition of $\langle \mathcal{L} \models \mathcal{R}; p \rangle$ to the set of assumed obligations $\tilde{A}$ is sound depends on the use of them in the application of (IA-1) and (IA-2).

We now proceed to prove the soundness of (IA-1) and (IA-2). First, define a *refutation* to an obligation $\mathcal{L} \models \mathcal{R}$ as a successful derivation of one or more atoms in $\mathcal{L}$ whose answer Ψ has an instance (ground substitution) β such that $\Psi\beta \wedge \mathcal{R}\beta$ is false. A finite refutation corresponds to a such derivation of finite length. A nonexistence of finite refutation means that $lm(P) \models (\mathcal{L} \rightarrow \mathcal{R})$, or in other words, $\mathcal{L} \models \mathcal{R}$. A derivation of a refutation is obtainable by left unfold (LU+I)) rule only. Hence a finite refutation of length k implies a corresponding k left unfold (LU+I) applications that results in a contradiction.

In the rules (IA-1) and (IA-2), we assume the hypothesis $A\theta$, where $A \equiv \mathcal{L}' \models \mathcal{R}'$ is some entailment encountered previously. By having the side conditions proved separately, we then can soundly transform the current entailment B into a new entailment C. In case of (IA-1), $B \equiv p(\tilde{x}) \wedge \mathcal{L}_1 \wedge \mathcal{L}_2 \models \mathcal{R}$ and $C \equiv \mathcal{R}'\theta \wedge \mathcal{L}_2 \models \mathcal{R}$. In case of (IA-2), we have $B \equiv p(\tilde{x}) \wedge \mathcal{L}_1 \models \mathcal{R}$ and $C \equiv \mathcal{L}_1 \models \mathcal{L}'\theta$.

Notice that the side conditions ensure that $A\theta \implies (C \implies B)$, where $\implies$ denotes implication. The side conditions also enforce the renaming θ to "progress" at least the left unfold of recursive atom $p(\tilde{y})$ to match with a newly generated atom $p(\tilde{x})$. This indeed enforces a well-founded measure on A.

To be more concrete, note that our transformation from B to C is *unsound only if* there exists a refutation β to B, and therefore A, but β is not a refutation to C. We then proceed to prove by contradiction. W.l.o.g., assume β is such a refutation and is the refutation to A which has the smallest length k. Trivially $k > 0$

as A has no finite refutation of length 0. Since there is at least one left unfold from A to B, β must be a refutation to B but of length less than equal to k. However, since $A\theta \implies (C \implies B)$, and β is a refutation of B but not C, therefore β is also a refutation of $A\theta$. Since θ must "progress" A by at least one left unfold, we end up with the fact that A has a refutation of length less than k. This is a contradiction. □

5. Implementation

```
function Prove(L ⊨ R, Ã, lb, rb, ib)
      /* Natural proof, i.e. by unification and SMT */
⟨1⟩   if (DirectProof(L ⊨ R)) return true

      let L = Φ, p_1, ..., p_n and R = Ψ, q_1, ..., q_m
⟨2⟩   OrSet := ∅
⟨3⟩   if (ib < INDUCTIONBOUND)      /* Induction Application */
⟨4⟩     foreach(⟨p(ỹ) ∧ L' ⊨ R'); p(ỹ)⟩ ∈ Ã)
⟨5⟩       Find p(x̃) ∈ L s.t. gen(p(x̃)) ≥ kill(p(ỹ))
⟨6⟩       Find L_1 ⊆ L s.t. θ := DirectProof(L_1 ⊨ L') ≠ ⊥
⟨7⟩       if(x̃ = ỹθ and θ is a valid renaming)
⟨8⟩         L_new := L \ {p(x̃)} \ L_1 ∪ R'θ
⟨9⟩         OrSet ∪= {⟨L_new ⊨ R, Ã, lb, rb, ib + 1⟩}
⟨10⟩    foreach(⟨p(ỹ) ∧ L' ⊨ R'); p(ỹ)⟩ ∈ Ã)
⟨11⟩      if (θ_1 := DirectProof(R' ⊨ R) = ⊥) continue
⟨12⟩      Find p(x̃) ∈ L s.t. gen(p(x̃)) ≥ kill(p(ỹ))
⟨13⟩      Find a valid renaming θ ⊇ θ_1 s.t. x̃ = ỹθ
⟨14⟩      OrSet ∪= {⟨L \ {p(x̃)} ⊨ L'θ, Ã, lb, rb, ib + 1⟩}

⟨15⟩  if (lb < MAXLEFTBOUND)      /* Left Unfold */
⟨16⟩    foreach (p_i ∈ L)
⟨17⟩      Obs := ∅
⟨18⟩      Ã' := Ã ∪ {⟨L ⊨ R; p_i⟩}
⟨19⟩      foreach (L_j ∈ ({L_1, L_2, ... L_l} := UNFOLD(p_i)))
⟨20⟩        ob := ⟨(L_j ∪ L \ {p_i}) ⊨ R, Ã', lb + 1, rb, ib⟩
⟨21⟩        if (trivially_true(ob)) continue
⟨22⟩        Obs := Obs ∪ {ob}
⟨23⟩      if (Obs = ∅) return true else OrSet ∪= {Obs}

⟨24⟩  if (rb < MAXRIGHTBOUND and ¬contradict(L ⊨ R))
        /* Right Unfold */
⟨25⟩    foreach (q_i ∈ R)
⟨26⟩      foreach( R_j ∈ {R_1, R_2, ... R_k} := UNFOLD(q_i))
⟨27⟩        ob = ⟨L ⊨ (R_j ∪ R \ {q_i}), Ã, lb, rb + 1, ib⟩
⟨28⟩        OrSet ∪= {{ob}}

⟨29⟩  if (OrSet = ∅) return false
⟨30⟩  OrSet := OrderByHeuristics(OrSet)
⟨31⟩  foreach (Obs ∈ OrSet)
⟨32⟩    if (ProveAll(Obs)) return true
⟨33⟩  return false
endfunction
```

Figure 13: The Main Algorithm

Let us briefly highlight our implementation, which intuitively follows from the proof rules in Sec. 4. The main algorithm is in Figure 14. In the figure, we use $X \cup= Y$ to denote $X := X \cup Y$.

We start off by calling the function Prove with the original obligation $\mathcal{L} \models \mathcal{R}$, the set of assumptions $\tilde{A}$ to be $\emptyset$, and all the counters lb, rb, ib to be 0. The counters lb, rb, ib are to keep track of, respectively, how many left unfolds, right unfolds, and inductions have been applied in this current path. These counters are to ensure that our algorithm terminates. In our experiments, the typical values for `INDUCTIONBOUND, MAXLEFTBOUND, MAXRIGHTBOUND` are respectively 3, 5, 5.

Typically an unoptimized proof obligation usually can be partitioned into a number of smaller and simpler proof obligations (e.g., by eliminating redundant terms and variables). This step can be implemented using any standard proof slicing technique and is not the focus of our discussion.

```
function ProveAll(Obs)
⟨34⟩  foreach (⟨L ⊨ R, Ã, lb, rb, ib⟩ ∈ Obs)
⟨35⟩    if (¬ Prove(L ⊨ R, Ã, lb, rb, ib)) return false;
endfunction

function DirectProof(L ⊨ R)
⟨36⟩  if (∃ q(x̃) ∈ R such that ∄ q(ỹ) ∈ L) return ⊥
⟨37⟩  L' := get_all_recursive(L)
⟨38⟩  R' := get_all_recursive(R)
⟨39⟩  Θ := {substitution θ | R'θ ⊆ L'}
⟨40⟩  if (Θ = ∅) return ⊥
⟨41⟩  foreach (θ ∈ Θ)
⟨42⟩    Φ := get_all_nonrecursive(L)
⟨43⟩    Ψ := get_all_nonrecursive(R)
⟨44⟩    θ' := bind_remaining_existential_variables(Ψ, Φ, θ)
        /* Extend θ to θ' by trying obvious bindings
           for remaining existential variables */
⟨45⟩    if (has_existential_variables(Ψ, Φ, θ')) continue
⟨46⟩    if (entailment(Φ, Ψθ')) return θ'
⟨47⟩  return ⊥
endfunction
```

Figure 14: Supporting Functions

Base Case: The function DirectProof acts as the base case of our recursive algorithm. For each proof obligation, we first attempt a *direct proof*, i.e., to discharge by applying rule (SUB) repetitively and then querying Z3 solver [12], after treating all recursive predicates in the LHS as uninterpreted, as in (CP)-rule.

Intuitively, this step succeeds if the proof obligation is simple "enough" such that a *proof by matching* can be achieved. We note here that, our proof rules in Section 4 allow other rules, e.g., (RU) rule in Example 2, to interleave with the (SUB) and (CP) rules. However, in our deterministic implementation, applications of (SUB) and (CP) rules are coupled together.

Let us examine the function DirectProof. If there is a recursive predicate q on the RHS, but not in the LHS, the function returns immediately, indicating failure with $\bot$. Otherwise, the function then proceeds by finding some (not exhaustive) substitutions Θ such that with each $\theta \in \Theta$, we can simultaneously remove all the recursive predicates on the RHS. This process will remove most existential variables on the RHS, since existential variables usually appear in some recursive predicates.

In case there remain some existential variables on the RHS, we attempt to bind them with the obvious candidates on the LHS (therefore extend θ to θ'). After this attempt, if the RHS contains no existential variables, we then call an SMT solver for entailment check. If the answer is yes, θ' is returned, indicating that a direct proof has been achieved.

Recursive Call: When the attempt of direct proof is not successful, we collect all possible transformations of the current proof obligation, using (IA-1), (IA-2), (LU+I), (RU) rules, into *a set of set of* obligations $OrSet$. The current proof obligation can be successfully discharged if there is *any* set of proof obligations $Obs \in OrSet$, where we can discharge *every* proof obligation $ob \in Obs$. The realization of the proof rules in our algorithm is straightforward, except for a few noteworthy points:

1. Our induction applications will not exhaustively search for all possible candidates. Instead, we only search for some trivial renaming which meets the side conditions of the rules.
2. When we perform left unfold, an obligation which is trivially true (trivially_true), i.e. the non-recursive part of the LHS is unsatisfiable, is immediately removed.
3. If the current obligation contains the LHS and RHS which contradict each other (contradict), right unfold will be avoided. The proof for this obligation can succeed only if there are no models for the LHS (so only left unfolds are required).

We note that our proof search proceeds recursively in a depth first search manner. The order in which the sets of obligations $Obs \in OrSet$ are considered might heavily affect the efficiency, i.e. the running time, but *not* the *effectiveness*, i.e. the ability to prove, of our framework. Such order is dictated by our heuristics, as the call to function OrderByHeuristics (line 30) indicates. We remark that our heuristics, described below, are very *intuitive* and directly follow from the fact that our base case is reached by a successful direct proof.

We proceed by a number of passes. In each pass, we first order the obligations within each $Obs \in OrderSet$. We then consider the order of $OrderSet$ by comparing the last obligation in each set $Obs \in OrderSet$. Subsequent passes will not undo the work of the previous passes, but instead work on the obligations and/or sets of obligations which are tied in previous passes.

1. An obligation which has contradicting LHS and RHS, given by the function contradict will be ordered after those do not (since the chance to successfully discharge such obligation is small).
2. An obligation contains no recursive predicates on the RHS will be order before those contain some recursive predicate(s) on the RHS.
3. An obligation having a recursive predicate q such that q appears in the RHS but not in the LHS will be ordered after those not.
4. An obligation contains more existential variables which cannot be deterministically bound to some non-existential variables (variables on the LHS) will be ordered after those contains less.
5. An obligation resulted from a left unfold will be ordered before those resulted from a right unfold (since it allows an induction hypothesis to be added).

EXAMPLE 5. *Revisit the obligation in Example 2, but now with the starting set of assumptions to be empty:*

$$\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2\,.$$

For simplicity we ignore the information about the counters lb, rb, and ib. First, the call to DirectProof fails since there is the predicate ls which appears in the RHS but not in the LHS. Induction rules cannot take place either, as the set of assumptions is currently empty. We prooceed by performing a left unfold first. Note that there is only one recursive predicate on the LHS. Let $\tilde{A}$ be:

$$\{\langle \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2; \mathsf{list}(x, L)\rangle\}.$$

The result for our left unfold is a set of two obligations:

$$O_0 \equiv \{\tilde{A}' \vdash x = 0, L \doteq \Omega \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2\,;\\ \tilde{A} \vdash L{\doteq}(x \mapsto t){*}L', \mathsf{list}(t, L') \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L{\doteq}L_1{*}L_2\}$$

We proceed with right unfold, producing four sets, each consists of one (simplified) obligation as follows:

$$O_1 \equiv \{\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), y = 0, L_2 \doteq \Omega, L{\doteq}L_1{*}L_2\}\\ O_2 \equiv \{\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(t, L_3), L{\doteq}L_1{*}(y \mapsto t){*}L_3\}\\ O_3 \equiv \{\emptyset \vdash \mathsf{list}(x, L) \models x = y, \mathsf{list}(y, L_2), L \doteq \Omega{*}L_2\}\\ O_4 \equiv \{\emptyset \vdash \mathsf{list}(x, L) \models x{\neq}y, \mathsf{ls}(x, t, L_3), \mathsf{list}(y, L_2), L{\doteq}(t \mapsto y){*}L_3{*}L_2\}$$

Assume that the initial order of those sets of obligations are as shown above. After the first two passes, the order between those sets is the same. The third pass, however, moves the singleton set O_3 to the first position. The fourth pass, on the other hand, moves O_1 to the second position. The fifth pass keep O_0 at the third position. The remaining two singleton sets, namely O_2 and O_4 are tied and placed at the end.

We proceed with the first obligation set, namely O_3, and a direct proof of it is successful. Therefore the original obligation can be discharged. The corresponding sequence of the proof rules is shown below, which is slightly different from what shown in Fig. 7.

$$\text{(RU)}\ \frac{\text{(SUB)}\ \dfrac{\text{(CP)}\ \dfrac{\texttt{True}}{\emptyset \vdash \mathsf{list}(x, L) \models x = x, L \doteq \Omega{*}L}}{\emptyset \vdash \mathsf{list}(x, L) \models x = y, \mathsf{list}(y, L_2), L \doteq \Omega{*}L_2}}{\emptyset \vdash \mathsf{list}(x, L) \models \mathsf{ls}(x, y, L_1), \mathsf{list}(y, L_2), L \doteq L_1{*}L_2}$$

6. Experiments

Our evaluations are performed on a 3.2Gz Intel processor with 2GB RAM, running Linux. We evaluated our prototype on a comprehensive set of benchmarks, including both academic algorithms and real programs. The benchmarks are collected from existing systems [27, 9, 23, 30, 7], those considered as the state-of-the-art for the purpose of proving user-defined recursive data-structure properties in imperative languages. We first demonstrate our evaluation with benchmarks that the state-of-the-art can handle, then with ones that are beyond their current supports.

6.1 Within the State-of-the-art

In this subsection, we consider the set of proof obligations where the state-of-the-art, i.e., U+M and "Cyclic Proof", are effective. Given that our proof method theoretically subsumes the state-of-the-art, the main purpose of this study is to evaluate the *efficiency* of our implementation against existing systems. To some extent, this exercise serves as a sanity check for our implementation.

We first start with proof obligations where U+M can automatically discharge without the help of user-defined lemmas. They are collected from the benchmarks of U+M frameworks [9, 23, 30]. As expected, our prototype prove all of those obligations, the running time for each is negligible (less than 0.2 second). This is because the proof obligations usually require just either *one* left unfold or *one* right unfold before matching — i.e., a direct proof — can successfully take place.

The second set of benchmarks are from "Cyclic Proof" [7], which are also used in SMT-COMP 2014 (Separation Logic)[6]. They are proof obligations which involve unmatchable predicates, thus U+M will not be effective. We also succeed in proving all of those obligations, less than a second for each.

In summary, the results demonstrate that (1) our prototype is able to handle what the state-of-the-art can; (2) our implementation is *competitive* enough.

6.2 Beyond the State-of-the-art

We now move to demonstrate the *key* result of this paper: proving what are beyond the state-of-the-art.

Proving User-Defined Lemmas: Our prototype can prove all commonly used lemmas, collected from [27, 9, 23, 30], which U+M and "Cyclic Proof" cannot handle. The running time is always less than a second for each lemma. Table 1 shows a non-exhaustive list of common user-defined lemmas. We purposely abstract them from the original usage in order to make them general and representative enough. The lemmas are written in traditional Separation Logic syntax for succinctness. Note that due to the duality of the definitions for list segments, e.g., ls vs. $\widehat{\mathsf{ls}}$, each lemma containing them would usually has a dual version, which for space reason we do not list down in Table 1. Similarly, some extensions, e.g., to capture the relationship of collective data values (using sets or sequences) between the LHS and the RHS, while can be automatically discharged by our prototype, are not listed in the table.

[6] See https://github.com/mihasighi/smtcomp14-sl

Lemma
$\mathsf{sorted\_list}(x, min) \models \mathsf{list}(x)$
$\mathsf{sorted\_list}_1(x, len, min) \models \mathsf{list}_1(x, len)$
$\mathsf{sorted\_list}_1(x, len, min) \models \mathsf{sorted\_list}(x, min)$
$\mathsf{sorted\_ls}(x, y, min, max) * \mathsf{sorted\_list}(y, min_2) \wedge max \leq min_2 \models \mathsf{sorted\_list}(x, min)$
$\mathsf{ls}(x, y) * \mathsf{list}(y) \models \mathsf{list}(x)$
$\mathsf{ls}(x, y) \models \widehat{\mathsf{ls}}(x, y)$ and $\widehat{\mathsf{ls}}(x, y) \models \mathsf{ls}(x, y)$
$\widehat{\mathsf{ls}}_1(x, y, len_1) * \widehat{\mathsf{ls}}_1(y, z, len_2) \models \widehat{\mathsf{ls}}_1(x, z, len_1{+}len_2)$
$\mathsf{ls}_1(x, y, len_1) * \mathsf{list}_1(y, len_2) \models \mathsf{list}_1(x, len_1{+}len_2)$
$\widehat{\mathsf{ls}}_1(x, last, len) * (last \mapsto new) \models \widehat{\mathsf{ls}}_1(x, new, len + 1)$
$\mathsf{dls}(x, y) * \mathsf{dlist}(y) \models \mathsf{dlist}(x)$
$\widehat{\mathsf{dls}}_1(x, y, len_1) * \widehat{\mathsf{dls}}_1(y, z, len_2) \models \widehat{\mathsf{dls}}_1(x, z, len_1{+}len_2)$
$\mathsf{dls}_1(x, y, len_1) * \mathsf{dlist}_1(y, len_2) \models \mathsf{dlist}_1(x, len_1{+}len_2)$
$\mathsf{avl}(x, height, min, max, balance) \models \mathsf{bstree}(x, height, min, max)$
$\mathsf{bstree}(x, height, min, max) \models \mathsf{bintree}(x, height)$

Table 1. Proving lemmas (existing systems cannot prove).

Let us briefly discuss Table 1. The first group talks about sorted linked lists. As an example, the second lemma is to state that a sorted list with length len and the minimum element min is also a list with the same length. We use indexes for different definitions of data structures, which involve different properties (i.e., sorted_list and $\mathsf{sorted\_list}_1$). The second, third and fourth groups are related to singly-linked lists, doubly-linked lists, and trees respectively.

Verifying Programs without Using Lemmas: Lemmas can serve many purposes. One important usage of lemmas in U+M systems is to equip a proof system with the power of user-provided re-writing rules, so as to overcome the main limitation of unfold-and-match. However, in the context of program verification, eliminating the usage of lemmas is crucial for improving the performance. This is because lemma applications, coupled with unfolding, often induce very large search space.

We now use the set of academic algorithms and open-source library programs[7], collected and published by [9, 30], to demonstrate that our prototype can verify all of the programs in this set without using lemmas. The library programs include Glib open source library, the OpenBSD library, the Linux kernel, the memory regions and the page cache implementations from two different operating systems. While Table 2 summarizes the verification of data structures from academic algorithms, Table 3 reports on open-source library programs.

Data-structure	Function	T/F
`Sorted List`	`find_last_iter, insert_iter, quick_sort_iter, bubble_sort`	$<1s$
`Circular List`	`count`	$<1s$
`Binary Search Tree`	`insert_iter,find_leftmost_iter, remove_root_iter, delete_iter`	$<1s$

Table 2. Verification of Academic Algorithms (existing systems require lemmas).

Remark #1: Using automatic induction, we have successfully eliminated the requirement for lemmas in existing systems (e.g. [9, 30]) for proving the functional correctness of the programs in Table 2 and 3. As already stated in Sec. 1, existing systems require lemmas in two common scenarios. First, it is when the traversal order of the data structures is different from what suggested by the recursive definitions, e.g. `OpenBSD/queue.h`. Second, it is due to the boundaries caused by iterative loops or multiple function calls. One example is `append` function in `glib/gslist.c`, where (in addition to the list definition) the list segment, ls(head,last), is necessary to say about the function invariant — the last node of a non-empty input list is always reachable from the list's head. Other examples are to make a connection between a sorted list and a singly-linked list (e.g. in sorting algorithms), between two sorted partitions (e.g. in `quick_sort_iter`), between a circular list and a list segment (e.g. `count`), etc.

[7] See `http://www.cs.uiuc.edu/~madhu/dryad/sl`

Program	Function	T/F
glib/gslist.c `Singly Linked-List`	`find, position, index, nth,last,length,append, insert_at_pos,merge_sort, remove,insert_sorted_list`	$<1s$
glib/glist.c `Doubly Linked-List`	`nth, position, find, index, last, length`	$<1s$
OpenBSD/queue.h `Queue`	`simpleq_remove_after, simpleq_insert_tail, simpleq_insert_after`	$<1s$
ExpressOS/cachePage.c	`lookup_prev,add_cachepage`	$<1s$
linux/mmap.c	`insert_vm_struct`	$<1s$

Table 3. Verification of Open-Source Libraries (existing systems require lemmas).

Remark #2: The verification time for each function is always less than 1 second. This is within our expectation since when our proof method succeeds, the size of the proof tree is relatively small. For example, in order to prove the functional correctness of `append` function in `glib/gslist.c`, we only need to prove 3 obligations, each of which requires no more than two left unfolds, two right unfolds and two inductions[8]. In fact, the maximum number of left unfolds, right unfolds and inductions used in our system are 5, 5 and 3 respectively, even for the functions that take U+M frameworks much longer time to prove. For example, consider `simpleq_insert_after`, a function to insert an element into a queue. This example requires reasoning about unmatchable predicates: to prove it DRYAD needs 18 seconds and the help from a lemma. Such inefficiency is due to the use of a complicated lemma[9], which consists of a large disjunction. Though efficient in practice, SMT solvers still face a combinatorial explosion challenge as they dissect the disjunction. In other words, in addition to having a higher level of automation, our framework has a potential advantage of being more efficient than existing U+M systems.

7. Related Work

There is a vast literature on program verification considering data structures. The well known formalism of Separation Logic (SL) [35] is often combined with a recursive formulation of data structure properties. Implementations, however, are incomplete, e.g., [2, 15], or deal only with fragments [1, 24]. There is also literature on decision procedures for restricted heap logics; we mention just a few examples: [31, 32, 19, 33, 4, 3]. These have, however, severe restrictions on expressivity. None of them can handle the VC's of the kind considered in this paper.

There is also a variety of verification tools based on classical logics and SMT solvers. Some examples are Dafny [21], VCC [10] and Verifast [16] which require significant ghost annotations, and annotations that explicitly express and manipulate frames. They do not automatically verify the general and complex properties addressed in this paper, but in general resort to interactive theorem

[8] Since the number of rules (disjuncts) in a predicate definition is fixed and small, the size of proof tree mainly depends on the number of unfolds and inductions.

[9] We believe that the lemmas in [30] are unnecessarily too complicated, because the authors want to reduce the number of them, by grouping a few into one.

provers, e.g. Mona, Isabelle or Coq, which usually requires manual guidance.

In [25], Navarro and Rybalchenko showed that significant performance improvements can be obtained by incorporating first-order theorem proving techniques into SL provers. However, the focus of that work is about list segments, not general user-defined recursive predicates. On a similar thread, [29] advances the automation of SL, using SMT, in verifying procedures manipulating list-like data structures. The works [39, 40, 9, 23, 30] are also close related works: they form the U+M paradigm which we have carefully discussed in Section 1 and 2.

We have discussed the works on automatic (and explicit) induction [5, 13, 22, 37] in Section 2. Here we further highlight the work of Lahiri and Qadeer [20], which adapts the induction principle for proving properties of well-founded linked list. The technique relies on the well-foundedness of the heap, while employing the induction principle to derive from two basic axioms a small set of additional first-order axioms that are useful for proving the correctness of several simple programs.

Closer to the spirit of our work, there are works on "Cyclic Proof" [6, 7] (which we have discussed) and "Matching Logic" [36]. They are based on the same principle that when two similar obligations are detected in the same proof path, the latter can be used to subsume to former. To some extent, these are special and limited instances of our induction rules. In terms of implementation, the "circularity rule" in [36] can only be applied to basic "patterns" in the logic, therefore it cannot support general user-defined recursive predicates.

We finally mention the work [18], from which the concept of our automatic induction originates. The current paper extends [18] first by refining the original single coinduction rule into two more powerful rules, to deal with the antecedent and consequent of a VC respectively. Secondly, the *application* of the rules has been systematized so as to produce a rigorous proof search strategy. Another technical advance is our introduction of *timestamps* (a progressive measure) in the two induction rules as an efficient technique to avoid circular reasoning. Finally, the present paper focuses on program verification and uses a specific domain of discourse involving the use of explicit symbolic heaps and separation.

8. Concluding Remarks

We presented a framework for proving recursive properties of data structures. The main contribution was an algorithm which provided a new level of automation across a wider class of programs. Its key technical features were two automatic re-writing rules, based on a systematic consideration of dynamically generated possibilities as induction hypotheses. Finally, experimental evidence showed that the algorithm has gone beyond the state-of-the-art.

References

[1] J. Berdine, C. Calcagno, and P. OHearn. A decidable fragment of separation logic. In *FSTTCS*, pages 97–109, 2004.

[2] J. Berdine, C. Calcagno, and P. OHearn. Symbolic execution with separation logic. In *APLAS*, pages 52–68, 2005.

[3] N. Bjrner and J. Hendrix. Linear functional fixed-points. In *CAV, LNCS 5643*, pages 124–139. Springer, 2009.

[4] A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. A logic-based framework for reasoning about composite data structures. In *CONCUR, LNCS 5710*, pages 178–195. Springer, 2009.

[5] R. S. Boyer and J. S. Moore. A theorem prover for a computational logic. In *CADE*, 1990.

[6] J. Brotherston, D. Distefano, and R. L. Petersen. Automated cyclic entailment proofs in separation logic. In *CADE*, 2011.

[7] J. Brotherston, N. Gorogiannis, and R. L. Petersen. A generic cyclic theorem prover. In *APLAS*, pages 350–367, 2012.

[8] J. Brotherston and J. Villard. Parametric completeness for separation theories. In *POPL*, pages 453–464, 2014.

[9] W.-N. Chin, C. David, H. H. Nguyen, and S. Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. In *SCP*, pages 1006–1036, 2012.

[10] E. Cohen, M. Dahlweid, M. A. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies. Vcc: A practical system for verifying concurrent c. In *TPHOLs*, 2009.

[11] B. Cook, C. Haase, J. Ouaknine, M. J. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In *CONCUR*, pages 235–249, 2011.

[12] L. De Moura and N. Bjørner. Z3: an efficient smt solver. In *TACAS*, 2008.

[13] P. C. Dillinger, P. Manolios, D. Vroon, and J. S. Moore. ACL2s: "The ACL2 Sedan". In *ICSE*, 2007.

[14] G. Duck, J. Jaffar, and N. Koh. A constraint solver for heaps with separation. In *CP, LNCS 8124*, 2013.

[15] R. Iosif and A. Rogalewicz abd J. Simachek. The tree width of separation logic with recursive definitions. In *CADE*, 2013.

[16] B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens. Verifast: A powerful, sound, predictable, fast verifier for c and java. In *NFM*, pages 41–55, 2011.

[17] J. Jaffar and M. J. Maher. Constraint logic programming: A survey. *J. LP*, 19/20:503–581, May/July 1994.

[18] J. Jaffar, A. E. Santosa, and R. Voicu. A coinduction rule for entailment of recursively defined properties. In *CP*, 2008.

[19] S. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using smt solvers. In *POPL*, pages 171–182. ACM, 2008.

[20] S. K. Lahiri and Shaz Qadeer. Verifying properties of well-founded linked lists. In *POPL*, pages 115–126. ACM, 2006.

[21] K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *LPAR, LNCS 6355*, pages 348–370. Springer, 2010.

[22] K. R. M. Leino. Automating induction with an smt solver. In *VMCAI*, 2012.

[23] P. Madhusudan, X. Qiu, and A. Stefanescu. Recursive proofs for inductive tree data-structures. In *POPL*, pages 611–622. ACM, 2012.

[24] S. Magill, M.-H. Tsai, P. Lee, and Y.-K. Tsay. Thor: A tool for reasoning about shape and arithmetic. In *CAV*, 2008.

[25] P. Navarro and A. Rybalchenko. Separation logic + superposition calculus = heap theorem prover. In *PLDI*, 2011.

[26] P. Navarro and A. Rybalchenko. Separation logic modulo theories. In *APLAS*, pages 90–106, 2013.

[27] H. H. Nguyen and W. N. Chin. Enhancing program verification with lemmas. In *CAV '08*, pages 355–369. Springer-Verlag, 2008.

[28] E. Pek, X. Qiu, and P. Madhusudan. Natural proofs for data structure manipulation in c using separation logic. In *PLDI*, 2014.

[29] R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using smt. In *CAV*, 2013.

[30] X. Qiu, P. Garg, A. Stefanescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In *PLDI*, pages 231–242, 2013.

[31] Z. Rakamaric, J. D. Bingham, and A. J. Hu. An inference-rule-based decision procedure for verification of heap-manipulating programs with mutable data and cyclic data structures. In *VMCAI*, 2007.

[32] Z. Rakamaric, R. Bruttomesso, A. J. Hu, and A. Cimatti. Verifying heap-manipulating programs in an smt framework. In *ATVA*, 2007.

[33] S. Ranise and C. Zarba. A theory of singly-linked lists and its extensible decision procedure. In *SEFM*, pages 206–215, 2006.

[34] J. Reynolds. A short course on separation logic. `http://www.cs.cmu.edu/afs/cs.cmu.edu/project/fox-19/member/jcr/wwwaac2003/aac.html`, 2003.

[35] J. C. Reynolds. Separation logic: A logic for shared mutable data objects. In *LICS*. IEEE, 2002.

[36] G. Rosu and A. Stefanescu. Checking reachability using matching logic. In *OOPSLA '12*, pages 555–574. ACM, 2012.

[37] W. Sonnex, S. Drossopoulou, and S. Eisenbach. Zeno: An automated prover for properties of recursive data structures. In *TACAS*, 2012.

[38] G. Stewart, L. Beringer, and A. W. Appel. Verified heap theorem prover by paramodulation. In *ICFP*, pages 3–14. ACM, 2012.

[39] K. Zee, V. Kuncak, and M. Rinard. Full functional verification of linked data structures. In *PLDI*, pages 349–361. ACM, 2008.

[40] K. Zee, V. Kuncak, and M. Rinard. An integrated proof language for imperative programs. In *PLDI*, pages 338–351. ACM, 2009.